

Oracle® Communications

Network Analytics Data Director

Benchmarking Guide



Release 25.2.100
G41331-01
October 2025

ORACLE®

G41331-01

Copyright © 2023, 2025, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

1	Introduction	
1.1	Purpose and Scope	1
1.2	References	1
2	Deployment	
3	Resource Requirement	
3.1	Profile Resource Requirements	2
3.1.1	Resource Profile for Database	2
3.1.2	Resource Profile for 500K MPS	3
3.1.2.1	Resource Profile for OCNADD OAM Services	3
3.1.2.2	Resource Profile for OCNADD Worker Group Services	4
3.2	Pod Affinity (or Anti-affinity) Rules	6
3.3	Ephemeral Storage Requirements	7
3.4	Disk Throughput Requirements	8
3.5	Kafka PVC Storage Requirements	10
4	OCNADD Benchmarking Testing	
4.1	Bare Metal Cluster Environment	1
4.1.1	Performance Benchmarking for 500K MPS SCP and 505K Egress Traffic with Two Synthetic Feeds with Two Endpoints	1
4.1.2	Performance Benchmarking for 360K MPS SCP and 450K Egress Traffic with Synthetic (TCP) and HTTP2 Feed	4
4.1.3	Performance Benchmarking for 360K MPS SCP Traffic with Synthetic (TCP) Feed	7
4.1.4	Performance Benchmarking for 270K MPS Aggregated Traffic with HTTP2 Feed	10
4.1.5	Performance Benchmarking for 270K MPS Aggregated Traffic with Synthetic (TCP) Feed	14
4.1.6	Performance Benchmarking for 270K MPS SCP Traffic with Kafka Feed	18
4.1.7	Performance Benchmarking for 270K MPS SCP Traffic with Synthetic Feed	20
4.1.8	Performance Benchmarking for 135K MPS Traffic with Synthetic Feed Replication	23

4.1.9	Performance Benchmarking for 30K MPS PCF Traffic with HTTP2 Feed	26
4.1.10	Performance Benchmarking for 9K BSF Traffic with HTTP2 Feed	29
4.2	vcNE Cluster Environment	31
4.2.1	Performance Benchmarking for 360K MPS SCP and 450K Egress Traffic with TCP and HTTP2 Feed	31
4.2.2	Performance Benchmarking for 360K MPS SCP Traffic with Synthetic (TCP) Feed	34

Preface

- [Documentation Accessibility](#)
- [Diversity and Inclusion](#)
- [Conventions](#)

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

My Oracle Support (MOS)

My Oracle Support (<https://support.oracle.com>) is your initial point of contact for all product support and training needs. A representative at Customer Access Support can assist you with My Oracle Support registration.

Call the Customer Access Support main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at <http://www.oracle.com/us/support/contact/index.html>. When calling, make the selections in the sequence shown below on the Support telephone menu:

- For Technical issues such as creating a new Service Request (SR), select **1**.
- For Non-technical issues such as registration or assistance with My Oracle Support, select **2**.
- For Hardware, Networking and Solaris Operating System Support, select **3**.

You are connected to a live agent who can assist you with My Oracle Support registration and opening a support ticket.

My Oracle Support is available 24 hours a day, 7 days a week, 365 days a year.

Acronyms

The following table provides information about the acronyms used in the document.

Table Acronyms

Acronym	Description
BSF	Oracle Communications Cloud Native Core, Binding Support Function
CNE	Oracle Communications Cloud Native Core, Cloud Native Environment
MPS	Messages Per Second
NDB	Network Data Broker
NRF	Oracle Communications Cloud Native Core, Network Repository Function
NVME	Non Volatile Memory Express
OCI	Oracle Cloud Infrastructure
OCNADD	Oracle Communications Network Analytics Data Director
OCPU	Oracle Compute Unit
PCF	Oracle Communications Cloud Native Core, Policy Control Function
SCP	Oracle Communications Cloud Native Core, Service Communication Proxy
SEPP	Oracle Communications Cloud Native Core, Security Edge Protection Proxy
xDR	Extended Detail Record
LB	Load Balancer
LBVM	Load Balancer Virtual Machine
CNLB	Cloud Native Load Balancer
PVC	Persistent Volume Claim

What's New in This Guide

This section lists the documentation updates for Release 25.2.1xx in *Oracle Communications Network Analytics Data Director Benchmarking Guide*.

Release 25.2.100 - G41331-01, October 2025

Updated the following sections:

- Updated [Resource Profile for OCNADD OAM Services](#).
- Added resource profile for [Egress 500K MPS Synthetic \(TCP\) Feed](#).
- Updated [Ephemeral Storage Requirements](#) table.
- Added the following scenarios in [OCNADD Benchmarking Testing](#) chapter:
 - [Performance Benchmarking for 500K MPS SCP and 505K Egress Traffic with Two Synthetic Feeds with Two Endpoints](#)

1

Introduction

Oracle Communications Network Analytics Data Director (OCNADD) is a specialized Network Data Broker (NDB) in 5G Network Architecture.

OCNADD receives network traffic data from various sources, 5G network functions (NFs), Non-5G NFs, and third-party producers, performs filtering, replication, and aggregation on the received data according to the rules implemented by the subscribed third-party consumers. OCNADD then sends the filtered, replicated, or aggregated data to the subscribed third-party consumers (third-party consumer applications or platforms) securely.

1.1 Purpose and Scope

This document is designed to measure the performance and capacity of OCNADD deployment and resource requirements.

1.2 References

For more information about OCNADD, see the following documents:


- *Oracle Communications Network Analytics Data Director Installation, Upgrade, and Fault Recovery Guide*
- *Oracle Communications Network Analytics Data Director User Guide*
- *Oracle Communications Cloud Native Core OCI Adaptor, NF Deployment on OCI Guide*

2

Deployment

OCNADD supports CNE and OCI deployment. There are a few variations in the deployment process for both the platforms. For more information about OCNADD installation, see *Oracle Communications Network Analytics Data Director Installation, Upgrade, and Fault Recovery Guide*.

The following diagram depicts the OCNADD deployment in the 5G architecture:



OCNADD uses the following common services of CNE:

- Kubernetes
- Prometheus
- Metallb (Load balancer)
- CNLB
- cnDBTier

3

Resource Requirement

This chapter provides information about the resource requirements to install and run Oracle Communications Network Analytics Data Director (OCNADD) with the desired Message Per Second (MPS) profiles.

Cluster Details

The following tables provides information about the types of servers and the number of servers used in the test environment:

Table 3-1 Test Bed 1 - CNE on BareMetal

Type of Server	X9 Server and NVME
Master node	3
Worker node	19
Storage Class	Standard
LB Type	LBVM/CNLB

Table 3-2 Test Bed 2 - vCNE on OpenStack

Type of Server	X9 Server and NVME
Master node	3
Worker node	44
Storage Class	Standard
LB Type	LBVM/CNLB

Resource Requirements for OCI Environment

- OCI block volume is attached to the PVC with auto-tune based performance from balanced to high performance. To change block volume to auto-tune based performance (Balance to High Performance), see [Changing the Performance of a Volume](#).
- All tests are performed with the default round-robin based ordering.
- Resource requirements may vary after enabling key or custom based ordering and running traffic with actual NFs.

Table 3-3 Test Bed 3 - OKE on OCI

Type of Server	OCI Hardware
Worker nodes	6
Instance Shape	VM.Standard.E4.Flex
OCPUs in worker node	50 (CPU: 100)
Memory in worker node	194 GB

3.1 Profile Resource Requirements

This section provides information about the profile resource requirements to install and run Oracle Communications Network Analytics Data Director (OCNADD) with the desired Message Per Second (MPS) profiles.

Note

It is recommended to have the following configurations for CNE Baremetal/vCNE setup to achieve the required throughput:

- Jumbo frames should be enabled.
- Ring buffer size should be increased to avoid packet drop at interfaces.(not applicable for vCNE)
- FluentD pods should not be in "CrashLoopBackOff" state due to Out of Memory error. For more information see "High Latency in adapter feeds due to high disk latency" section in *Oracle Communications Network Analytics Data Director Troubleshooting Guide*.
- The benchmark tests were performed with round trip latency of up to 5ms from third-party consumer applications. In case the latency is more than 5ms then the resource profile footprint and the E2E latency will be higher.

- [Resource Profile for Database](#)
- [Resource Profile for OCNADD OAM Services](#)
- [Resource Profile for OCNADD Worker Group Services](#)

3.1.1 Resource Profile for Database

This section provides information about the database profile resource requirements to install and run Oracle Communications Network Analytics Data Director (OCNADD) with the desired Message Per Second (MPS) profiles.

Table 3-4 Resource Requirement

cnDBTier Pods	Min vCPU	Max vCPU	Min Memory	Max Memory	Total Replica
SQL (ndbmysqld) Kubernetes Resource Type: StatefulSet	1	1	1Gi	1Gi	2
SQL (ndbappmysqld) Kubernetes Resource Type: Statefulset	1	1	1Gi	1Gi	2
MGMT (ndbmngmd) Kubernetes Resource Type: StatefulSet	1	1	1Gi	1Gi	2
Database (ndbmtd) Kubernetes Resource Type: StatefulSet	1	1	4Gi	4Gi	2

Table 3-4 (Cont.) Resource Requirement

cnDBTier Pods	Min vCPU	Max vCPU	Min Memory	Max Memory	Total Replica
Backup Manager Service (db-backup-manager-svc) Kubernetes Resource Type: Deployment	0.1	0.1	128Mi	128Mi	1
Monitor Service (db-monitor-svc) Kubernetes Resource Type: Deployment	0.2	0.2	500Mi	500Mi	1
EXTENDED STORAGE is ENABLED in CORRELATION Feed(Per Correlation Feed) Rate Supported in current release: 1K MPS rate with 24 hours retention Update "global.ndb.datamemory=96G" in custom-value.yaml of cndbTier PVC of ndbmt= 150GB					
Database (ndbmt) Kubernetes Resource Type: StatefulSet	8	8	128Gi	128Gi	4

Note

Configure "**datamemory: 1G**" under "**ndbmt**" section while deploying the CnDbTier for OCNADD. For more details on cnDBTier resource profile, see "cnDBTier Small Profile" section in *cnDBTier Resource Models Guide*.

3.1.2 Resource Profile for 500K MPS

3.1.2.1 Resource Profile for OCNADD OAM Services

The following profile is used for management group services in all the performance scenarios.

Table 3-5 Resource Requirement

Service Name	vCPU	Memory Required (Gi)	Total Replica	Description
ocnaddconfiguration	1	1	1	-
ocnaddalarm	1	1	1	-
ocnaddhealthmonitoring	1	1	1	-
ocnaddgui	1	1	1	-
ocnadduirouter	1	1	1	-
ocnaddexport	0.5	1	1	Resource requirement will increase when export is configured.
ocnaddredundancyagent	1	1	1	Required only when Georedundancy is enabled for OCNADD.

3.1.2.2 Resource Profile for OCNADD Worker Group Services

The following profile shall be used for worker group services. The resource profile for worker group services will vary based on the scenario to be executed.

Note

To support the increased throughput, first the number of Kafka instances must be increased followed by the number of topic partition changes based on the recommended MPS profile. For more details on this, see "Adding Partitions to an Existing Topic" in the *Oracle Communications Network Analytics Data Director User Guide*.

Note

For more details about various supported MPS resource profiles, refer to the *Oracle Communications Network Analytics Data Director Benchmarking Guide* of previous or respective releases.

3.1.2.2.1 Egress 500K MPS Synthetic (TCP) Feed

Note

Test conducted using separate clusters: One is dedicated to OCNADD services, while the other is shared by SCP NF services and third-party consumer.

- **Replication Factor:** 1
- **Message Size:** 3500
- **Feed Type:** 2-TCP feeds
- **FILTER:** Egress filter ON for second Feed with 1% data
- **Message Sequencing/Metadata:** OFF
- **Test bed:** LBVM-CNE Bare Metal Cluster Environment
- **Storage Type:** Ceph

500K MPS SCP Profile(Ingress to DD)

Table 3-6 500K MPS SCP Profile (Ingress to DD)

Service	vCPU	Memory Required (Gi)	Total Replica	Topic Partitions
kraft-controller	1	2	3	-
ocnaddkafka (kafkaBroker)	5	96	17	-
ocnaddscpaggregation	2	4	20	SCP =120 (Each instance 6 partition)

Table 3-6 (Cont.) 500K MPS SCP Profile (Ingress to DD)

Service	vCPU	Memory Required (Gi)	Total Replica	Topic Partitions
ocnaddadapter-1(T CP) (consumeradapter, without any filter)	3	6	28	MAIN=168 (Each instance 6 partition)
ocnaddadapter-2(T CP) (consumeradapter, 500K MPS ingress and 5K MPS egress with Filter)	3	10	9	No partition change as MAIN topic already has 168 partition.
ocnaddadminservice	1	1	1	Admin Service moved to Worker-Group from release 25.2.100

Note

- Additional memory and/or replicas are required for the aggregation service if the **Metadata Cache** feature is enabled and the values of the properties `METADATA_MAP_CACHE_EXPIRY_TIME_MS` and `METADATA_MAP_CACHE_SCHEDULER_TIME_MS` are increased to higher values.
- The end-to-end latency may increase based on:
 - Higher values of `METADATA_MAP_CACHE_EXPIRY_TIME_MS` and `METADATA_MAP_CACHE_SCHEDULER_TIME_MS`.
 - Timer Expiry Value + Processing Time + RF2/RF1 Processing Time + Third-party Response Time (for HTTP2 feed).
- Resource requirements may vary for the **consumeradapter service** based on the percentage of data allowed after filtering and the number of filter conditions along with their values.
- If **ramDriveStorage** is enabled, then the KafkaBroker Pod memory requirement will be: Memory Required for KafkaBroker + Memory Required for Data Retention
 Example: If the KafkaBroker requires **48Gi** of memory and data retention for the topic requires **200Gi**, the total memory required will be **248Gi**.
 It is recommended to configure a replication factor of **2** for the Kafka internal topics (`offsetTopicReplicationFactor` and `transactionStateReplicationFactor`) to improve cluster durability. This configuration may enhance data availability and resilience.
- Depending on cluster performance, more instances of the KafkaBroker may be required when running with **RF=2**, and end-to-end latency may also increase if disk I/O is slow.
- For DISK I/O, see [Disk Throughput Requirements](#).
- For Kafka PVC-Storage, see [Kafka PVC Storage Requirements](#).

3.2 Pod Affinity (or Anti-affinity) Rules

In the Data Director, support for node affinity has been added. The rules are currently defined for the services mentioned in the table below. The rules are currently disabled; however, the user can enable them for the supported services. The rules are provided to control the deployment of certain traffic processing services on a particular set of identified nodes.

Services name:

- Consumer Adapter
- Kafka
- ocnaddnrfaggregation
- ocnaddscpaggregation
- ocnaddseppaggregation
- ocnaddnonoracleaggregation
- ocnaddbsfaggregation
- ocnaddpcfaggregation

Node Affinity Rules

1. Update the "affinity" section in the ocnadd-custom-values.yaml file

```
affinity: {}

# Node Affinity Configuration:
#
# To enable node affinity, remove the empty curly braces
# ({{}}) above and un-comment the nodeAffinity section below.
# This allows you to specify rules for scheduling pods on
# specific nodes.
#
# Example Configuration:
#####
# nodeAffinity:
#   requiredDuringSchedulingIgnoredDuringExecution:
#     nodeSelectorTerms:
#       - matchExpressions:
#         - key: kubernetes.io/hostname
#           operator: NotIn
#           values:
#             - k8s-node-26
#             - k8s-node-24
#       - key: kubernetes.io/hostname
#         operator: In
#         values:
#           - k8s-node-2
#           - k8s-node-3
#####
# Explanation:
#
# - The 'NotIn' expression prevents pods from being
# scheduled on nodes k8s-node-26 and k8s-node-24.
# - The 'In' expression ensures pods are scheduled on
```



```

nodes k8s-node-2 and k8s-node-3.
#
# To customize, modify the 'key', 'operator', and 'values'
fields according to your needs.
# You can add or remove 'matchExpressions' to create more
complex scheduling rules.
#
# Remember to remove the empty 'affinity: {}' and un-
comment the desired nodeAffinity configuration to enable it.

```

2. Helm upgrade the corresponding worker group or the default group

```

helm upgrade <source-release-name> -f ocnadd-custom-values-<worker-
group>.yaml --namespace <source-release-namespace> <target-release-helm-
chart>

```

Where,

- <source-release-name> is the release name of the source release deployment.
- ocnadd-custom-values-<worker-group>.yaml is the custom values file created for default-worker-group or the Worker Group in separate namespace.
- <source-release-namespace> is the OCNADD namespace of the source release.
- <target-release-helm-chart> is the location of the Helm chart of the target release.

For example:

```

helm upgrade ocnadd -f ocnadd-custom-values-wg.yaml --namespace ocnadd-
deploy ocnadd_wg

```

3. Verify that PODs of the modified services have been deployed as per the configured affinity rules.

3.3 Ephemeral Storage Requirements

The following table describes the Ephemeral Storage requirements for OCNADD:

Table 3-7 Ephemeral Storage Requirements

Service Name	Ephemeral Storage (Request) in Mi	Ephemeral Storage (Limit) in Mi	Description
OAM Services			
ocnaddalarm	100	500	-
ocnaddhealthmonitoring	100	500	-
ocnaddconfiguration	100	500	-
ocnadduirouter	500	500	-
ocnaddexport	1000	2000	-
ocnaddredundancyagent	100	500	Required only when Geo Redundancy is enabled for OCNADD
Worker Group Services			

Table 3-7 (Cont.) Ephemeral Storage Requirements

Service Name	Ephemeral Storage (Request) in Mi	Ephemeral Storage (Limit) in Mi	Description
<app-name>-adapter (consumeradapter)	1000	1000	-
ocnaddscppaggregation	100	500	-
ocnaddseppaggregation	100	500	-
ocnaddnrffaggregation	100	500	-
ocnaddbsffaggregation	100	500	-
ocnaddpcfaggregation	100	500	-
ocnaddnonoracleaggregation	100	500	Required only when Data processing is enabled from Non-oracle NFs
ocnaddcorrelation	400	800	-
ocnaddstorageadapter	400	800	-
ocnaddingressadapter	400	800	-
ocnaddfilter	100	800	Required only when "Filtered" or "Correlated Filtered" feed is created
ocnaddadminservice	200	200	Admin Service moved to Worker-Group from release 25.2.100

3.4 Disk Throughput Requirements

The following table describes the disk throughput requirements in OCNADD:

Table 3-8 Disk Throughput Requirements

Avg Size (in Bytes)	Rate	RF (Kafka Replication Factor)	Topic (NF+M AIN)	Consumer Feed	Total Write Throughput (MB/s)	Total Read Throughput (MB/s)	No. of Broker	Per Broker Write Throughput (MB/s)	Per Broker Read Throughput (MB/s)	Total per Broker Throughput (MB/s) with 10% buffer	Total Disk Throughput (MB/s) for the Cluster with 10% Buffer
1941	39000	1	2	1	145	145	3	54	54	108	324
1941	39000	2	2	1	289	289	3	106	106	212	636
3769	39000	1	2	1	281	281	3	104	104	208	624
3769	39000	2	2	1	561	561	3	206	206	412	1236

Note

- The average size of OCNADD Ingress message captured in the table includes the size of metadata list + header list of original 5G HTTP2 header frame + 5G-SBI-Message.
- Currently, it is recommended to set the Replication Factor (RF) value to **1** with the assumption that the underlying storage provides data redundancy. RF value of "2" will be supported in a future release.

The disk throughput calculations are as follows:

Writes: $W * RF * T$
 Reads: $((RF*T)+C-1) * W$
 Disk Throughput (Write + Read): $(W * RF * T) + (L * W)$
 W -> MB/sec of data that will be written
 RF -> Replication factor
 T -> No of topics to which data copied. As of now, each message will be copied into two topics.
 C -> Number of consumer groups, that is the number of readers for each write
 L -> $(RF*T) + C - 1$

Average Message in Table:

Average Message Size = $(a_1b_1+a_2b_2+...+a(n)b(n))/(a_1+a_2+...+a(n))$
 a1 -> SCP MPS
 b1 -> SCP message size
 a2 -> NRF MPS
 b2 -> NRF message size
 a(n) -> NF(n) MPS
 b(n) -> NF(n) message size

Example:

Average message size for row 1 = $((1624*30000)+(3000*9000))/(30000+9000) = 1941$ Bytes (approx)

Average message size for row 4 = $((4000*30000)+(3000*9000))/(30000+9000) = 3769$ Bytes (approx)

The following table describes the disk throughput for SCP and NRF:

Table 3-9 SCP, NRF, and SEPP Disk Throughput

SCP Message		NRF Message	SEPP Message		RF (Kafka Replication Factor)	Topic (NF+MAIN)	Consumer Feed	Total Write Throughput (MB/s)	Total Read Throughput (MB/s)	No. of Broker	Per Broker Write Throughput (MB/s)	Per Broker Read Throughput (MB/s)	Total per Broker Throughput (MB/s) with 10% buffer	Total Disk Throughput (MB/s) for Cluster with 10% Buffer	Rate
Avg Size (Bytes)	Rate	Avg Size (Bytes)	Rate	Avg Size (Bytes)											
1624	30000	3000	9000	3000	15000	1	2	1	145	145	3	54	54	108	324
1624	30000	3000	9000	3000	15000	2	2	1	289	289	3	106	106	212	636
4000	30000	3000	9000	3000	15000	1	2	1	281	281	3	104	104	208	624
4000	30000	3000	9000	3000	15000	2	2	1	561	561	3	206	206	412	1236

Note

- The average size of OCNADD Ingress message captured in the table includes the size of metadata list + header list of original 5G HTTP2 header frame + 5G-SBI-Message.
- Currently, it is recommended to set the Replication Factor (RF) value to **1** with the assumption that the underlying storage provides data redundancy.

3.5 Kafka PVC Storage Requirements

The following table describes the retention period per topic for different NFs:

Table 3-10 Retention Period Per Topic

Topic Name	Retention Period
SCP	5 Minutes
NRF	5 Minutes
SEPP	5 Minutes
BSF	5 Minutes
PCF	5 Minutes
MAIN	6 Hours (Max)

The following calculation is for storage requirements for a topic:

! Important

For the 6 hrs storage in the MAIN topic, the storage requirement must be calculated using the following information:

Storage Requirement for a topic = MPS * Retention Period * RF * Average Message Size

Where,

MPS is "Message Per Second"

RF is "Replication Factor"

Examples:**1. Average Message Size = 1941 Bytes**

The following example uses the values from the first row of the [Table 3-9](#) table. For more information about the table, see [Disk Throughput Requirements](#):

Storage Requirement for SCP and NRF Topics = MPS * Retention Period * RF * Message Size
= 39000 * 5 Minutes * 3 * 1941

= 39000 * 5 * 60 * 3 * 1941

= ~ 63.45 GB

Storage Requirement for MAIN = MPS * Retention Period * RF * Message Size
= 39000 * 6 Hours * 3 * 1941

= 39000 * 6 * 60 * 60 * 3 * 1941

= ~ 4.46 TB

Total Storage Requirement for the Broker Cluster = Storage for SCP + Storage for NRF + Storage for MAIN

= 63.45 GB + 4.46 TB

= ~ 4.53 TB

Total Storage for each broker = (4.53/Number of Brokers) TB = (4.53/3) TB
= ~ 1.51 TB [Assuming 3 Broker cluster]

2. Average Message Size = 3769 Bytes

The following example uses the values from the fourth row of the [Table 3-9](#) table. For more information about the table, see [Disk Throughput Requirements](#):

Storage Requirement for SCP and NRF Topics = MPS * Retention Period * RF * Message Size
= 39000 * 5 Minutes * 3 * 3769

= 39000 * 5 * 60 * 3 * 3769

= ~ 123.20 GB

Storage Requirement for MAIN = MPS * Retention Period * RF * Message Size
= 39000 * 6 Hours * 3 * 3769

$$= 39000 * 6 * 60 * 60 * 3 * 3769$$

$$= \sim 8.66 \text{ TB}$$

Total Storage Requirement for the Broker Cluster = Storage for SCP +
Storage for NRF + Storage for MAIN

$$= 123.20 \text{ GB} + 8.66 \text{ TB}$$

$$= \sim 8.79 \text{ TB}$$

Total Storage for each broker = (8.79/Number of Brokers) TB = (8.79/3) TB
= ~ 2.93 TB [Assuming 3 Broker cluster]

4

OCNADD Benchmarking Testing

This section describes the performance testing scenarios and results for the Message Feed functionality provided by Oracle Communications Network Analytics Data Director (OCNADD). The message feed feature is tested with SCP, NRF, SEPP, BSF, and PCF as the source of the message feed.

4.1 Bare Metal Cluster Environment

This section gives the details of the performance tests done in Bare Metal Cluster Environment.

4.1.1 Performance Benchmarking for 500K MPS SCP and 505K Egress Traffic with Two Synthetic Feeds with Two Endpoints

This performance benchmarking evaluates the deployment of OCNADD handling an SCP traffic of 500K MPS and Egress traffic of 505K MPS. The benchmark utilizes two synthetic (TCP) feeds, with feed1(without filter) 500K MPS, Feed2(with 1-2% traffic filter) 5K-10K TCP Feeds:

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

- **Execution Time:** 24+ hours
- **Call Rate:** 500K MPS
- **Call Mix:** SCP Model-D

Setup Details

- **Environment:** OCCNE 25.1.200
- **LB Type:** Bare metal CNLB
- **Cluster Topology:**
 - **DD:** Deployed in Cluster-1
 - **NF:** Deployed in Cluster-2
 - **Third-Party Consumers:** Deployed in Cluster-3
- **Component Versions:**
 - **OCNADD:** 25.2.100
 - **cnDBTier:** 25.1.200

- **SCP Version:** 25.1.200
- **Third-Party Consumers:** Two endpoints with two different egress NADs (configured in round-robin)

Configuration Details

- **SCP:** messageCopy enabled
- **Security:** SASL/SSL enabled between SCP and OCNADD
- **OCNADD:**
 - Replication Factor: 1
 - Kafka PVC Size: 500 GB
- **Data Feed Configuration:**
 - **TCP Feed (1):** Ingress 500K, Egress 500K
 - **TCP Filtered Feed (2):** Ingress 500K, Egress 5K–10K
- **Message Size:** Approximately 3,500–4,000 bytes

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with SCP NF.
 - **Message Ingestion Rate:** 500K MPS

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with SCP.
 - Message Ingestion Rate: 500K MPS

Resource Specifications:

Table 4-1 Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddscppaggregation	2	2	4	4	20	20	120	SCP
ocnaddadapter-1(TCP-feed1)	3	3	6	6	28	28	168	MAIN

Table 4-1 (Cont.) Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddadapter-1(TCP-feed2)	3	3	10	10	9	9	168	MAIN
ocnaddkafka	5	5	96	96	17	17	-	-
kraft-controller	1	1	2	2	3	3	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnaddrouter	1	2	1	1	1	2	-	-

- **SCP:** SCP Traffic is 250K TPS using two trigger points.

Benchmark Test Results

Traffic Feed Details

Table 4-2 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	500K MPS	14	tcp feed1 51.4 ms tcp feed2 23.8 ms	99.98%
SCP	Ingress Gateway (Request) Egress Gateway (Request)	500K TPS	14	NA	NA

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-3 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replicas
ocnaddscpaggregation	82.1	27.16	20
ocnaddadapter-1(TCP-feed1)	82.1	27.8	28
ocnaddadapter-2(TCP-feed2)	85.3	29.5	9
kafka-broker	86.5	49	17
kraft-controller	0.65	55.75	3
ocnaddadminservice	0.10	24.80	1
ocnaddalarm	0.20	47.50	1
ocnaddconfiguration	0.20	32.40	1
ocnaddgui	0.10	4.90	1
ocnaddhealthmonitoring	0.30	33.30	1
ocnadduirouter	0.10	15.70	1

4.1.2 Performance Benchmarking for 360K MPS SCP and 450K Egress Traffic with Synthetic (TCP) and HTTP2 Feed

This performance benchmarking evaluates the deployment of OCNADD handling an SCP traffic of 360K MPS and Egress traffic of 450K MPS. The benchmark utilizes 3 synthetic (TCP) feeds and 1 HTTP2 feed:

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

- **Execution Time:** 24+ Hours
- **Call Rate:** 360K MPS
- **Call Mix:** SCP Model-C

Setup Details

- **Environment:** OCCNE 24.3.1

- **LB Type:** Bare metal CNLB
- **Cluster Topology:**
 - **DD:** Deployed in Cluster-1
 - **NF:** Deployed in Cluster-2
 - **Third-Party Consumers:** Deployed in Cluster-3
- **Component Versions:**
 - **OCNADD:** 25.1.200
 - **cnDBTier:** 25.1.100
 - **SCP Version:** 25.1.100
 - **Third-Party Consumers:** Two endpoints with two different egress NADs (configured in round-robin)

Configuration Details

- **SCP:** messageCopy enabled
- **Security:** SASL/SSL enabled between SCP and OCNADD
- **OCNADD:**
 - **Replication Factor:** 1
 - **Kafka PVC Size:** 400 GB
- **Data Feed Configuration:**
 - **TCP Feed:** Ingress 360K and Egress 360K
 - **TCP Filtered Feed (1):** Ingress 360K and Egress 36K
 - **TCP Filtered Feed (2):** Ingress 360K and Egress 36K
 - **HTTP Filtered Feed:** Ingress 360K and Egress 16K
- **Message Size:** Approximately 3500–4000 bytes

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with SCP.
 - Message Ingestion Rate: 360K MPS

Resource Specifications:

Table 4-4 Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-

Table 4-4 (Cont.) Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddscppaggregation	2	2	4	4	13	13	78	SCP
tcpfeedcnlb-sig1-sig2-egr1-adapter	3	3	6	6	20	20	120	MAIN
tcpfeedcnlb-2-sig3-sig4-egr1-adapter	3	3	10	10	7	7	120	MAIN
tcpfeedcnlb-sig3-sig4-egr1-adapter	3	3	10	10	7	7	120	MAIN
http2feed-sig3-sig4-egr1-cnlb-adapter	3	3	28	28	5	5	120	MAIN
ocnaddkafka	5	5	96	96	16	16	-	-
kraft-controller	1	1	2	2	3	3	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnadduirouter	1	2	1	1	1	2	-	-
ocnadduirouter	1	2	1	1	1	2	-	-
ocnaddfilter	2	2	3	3	1	1	-	-

- **SCP:** SCP Traffic is 180K TPS using two trigger points.

Benchmark Test Results

Traffic Feed Details

Table 4-5 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	360K MPS	24	httpfeed~ 21.6ms tcp feed1~ 22.2ms tcp feed3~ 28.5ms tcp feed3~ 22.7ms	NA

Table 4-5 (Cont.) Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
SCP	Ingress Gateway (Request) Egress Gateway (Request)	180K TPS	24	NA	NA

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-6 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
ocnaddscpaggregation	67.6	24.86	13
tcpfeedcnlb-sig1-sig2-egr1-adapter	65.55	36.25	20
tcpfeedcnlb-sig3-sig4-egr1-adapter	56.8	46.72	7
tcpfeedcnlb-2-sig3-sig4-egr1-adapter	56.85	46.72	7
http2feed-sig3-sig4-egr1-cnlb-adapter	65.6	55.52	5
kafka-broker	48.57	46.4	16
kraft-controller	0.5	47	3
ocnaddadminservice	0.1	36.8	1
ocnaddalarm	0.2	45.1	1
ocnaddconfiguration	0.2	47.6	1
ocnaddgui	0.1	3.3	1
ocnaddhealthmonitoring	0.2	38.8	1
ocnadduirouter	0.1	34.8	1

4.1.3 Performance Benchmarking for 360K MPS SCP Traffic with Synthetic (TCP) Feed

This performance benchmarking evaluates the deployment of OCNADD handling an SCP traffic of 360K MPS. The benchmark utilizes a synthetic (TCP) feed:

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

- **Execution Time:** 48+ Hours
- **Call Rate:** 360K MPS
- **Call Mix:** SCP Model-C

Setup Details

- **Environment:** OCCNE 24.3.1
- **LB Type:** Bare metal CNLB
- **Cluster Topology:**
 - **DD:** Deployed in Cluster-1
 - **NF:** Deployed in Cluster-2
 - **Third-Party Consumers:** Deployed in Cluster-3
- **Component Versions:**
 - **OCNADD:** 25.1.200
 - **cnDBTier:** 25.1.100
 - **SCP Version:** 25.1.100
 - **Third-Party Consumers:** Two endpoints with two different egress NADs (configured in round-robin)

Configuration Details

- **SCP:** messageCopy enabled
- **Security:** SASL/SSL enabled between SCP and OCNADD
- **OCNADD:**
 - **Replication Factor:** 1
 - **Data Feed Protocol:** Synthetic (TCP)
 - **Kafka PVC Size:** 400 GB
- **Message Size:** Approximately 3500–4000 bytes

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with SCP.
 - Message Ingestion Rate: 360 K MPS

Resource Specifications:

Table 4-7 Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Memory Request Per Pod (Gi)	Min Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddscppaggregation	2	2	4	4	13	13	78	SCP
ocnaddadapter	3	3	6	6	20	20	120	MAIN
ocnaddkafka	5	5	48	48	16	16	-	-
kraft-controller	1	1	2	2	3	3	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnadduirouter	1	2	1	1	1	2	-	-

- **SCP:** SCP Traffic is 180K TPS using two trigger points.

Benchmark Test Results

Traffic Feed Details

Table 4-8 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	360K MPS	48	19ms	NA
SCP	Ingress Gateway (Request) Egress Gateway (Request)	180K TPS	48	NA	NA

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-9 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
ocnaddscppaggregation	57.50	18.75	13
adapter	55.66	26.25	20
kafka-broker	35.15	12.8	16

Table 4-9 (Cont.) OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
kraft-controller	0.20	46.70	3
ocnaddadminservice	0.20	36.80	1
ocnaddalarm	0.20	45.10	1
ocnaddconfiguration	0.10	47.60	1
ocnaddgui	0.10	3.30	1
ocnaddhealthmonitoring	0.20	38.80	1
ocnadduirouter	0.10	34.80	1

4.1.4 Performance Benchmarking for 270K MPS Aggregated Traffic with HTTP2 Feed

This performance benchmarking evaluates the deployment of OCNADD handling an aggregated feed of 270K MPS. The breakdown includes 128K MPS SCP traffic (64K TPS with 2 trigger points), NRF 15K MPS (7.5K TPS with 2 trigger points) & SEPP 60K MPS (15K TPS with 4 trigger points), BSF 10K MPS (with 2 trigger points) PCF 58K MPS (with 2 trigger points). The benchmark utilizes an HTTP2 feed:

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

Execution Details

- **Execution Time:** 12+ Hours

Call Rate

- **SCP Traffic:** 128K MPS (64K TPS with 2 trigger points)
- **NRF:** 15K MPS (7.5K TPS with 2 trigger points)
- **SEPP:** 60K MPS (15K TPS with 4 trigger points)
- **BSF:** 10K MPS (2 trigger points)
- **PCF:** 58K MPS (2 trigger points)

Call Mix

- **NF Call Mix:** SCP, NRF, SEPP, BSF and PCF NF Call mix models

Setup Details

- **Environment:** OCCNE 24.3.1
- **LB Type:** Bare metal CNLB
- **Cluster Topology:**

- **DD:** Cluster-1
- **NF:** Cluster-2
- **Third-Party Consumers:** Cluster-3
- **OCNADD Version:** 25.1.200
- **cnDBTier:** 25.1.100

Network Functions (NFs)

- **SCP:** 25.1.100
- **NRF:** 25.1.100
- **SEPP:** 25.1.100
- **BSF:** 25.1.100
- **PCF:** 25.1.100

Third-Party Consumers

- Two endpoints with two different egress NADs
- **Load Balancing:** Weighted (50–50% or 60–40% distribution)

Configuration Details

- **NRF:** Ingress Gateway & Egress Gateway `messageCopy` enabled
- **SCP:** Default configuration
- **cSEPP:** Ingress PLMN/N32 and Egress PLMN/N32 `messageCopy` enabled
- **pSEPP:** `messageCopy` disabled
- **PCF:** Ingress Gateway & Egress Gateway `messageCopy` enabled
- **BSF:** Ingress Gateway & Egress Gateway `messageCopy` enabled

Security & Data Feed

- **SASL/SSL:** Enabled between BSF and OCNADD
- **OCNADD Replication Factor:** 1
- **OCNADD Data Feed:** HTTP2
- **OCNADD Kafka PVC:** 400 GB
- **Average Message Size:** ~3500 Bytes

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with aggregation rules set for SCP, NRF, SEPP, BSF, and PCF.
 - Message Ingestion Rate: 270K MPS

Resource Specifications:

Table 4-10 OCNADD Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Memory Request Per Pod (Gi)	Memory Limit Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddscppaggregation	2	2	4	4	5	5	30	SCP
ocnaddnrfaggregation	2	2	2	2	1	1	6	NRF
ocnaddseppaggregation	2	2	4	4	3	3	18	SEPP
ocnaddadapter	3	3	6	6	23	23	207	MAIN
ocnaddkafka	4	4	48	48	11	11	-	-
kraft-controller	1	1	2	2	3	3	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnadduirouter	1	2	1	1	1	2	-	-
ocnaddfilter	2	2	3	3	1	1	-	-
ocnaddexport	2	4	4	64	1	2	-	-
ocnaddpcfaggregation	2	2	2	2	2	2	12	PCF
ocnaddbsfaggregation	2	2	2	2	1	1	6	BSF

- **SCP Traffic:** 128K MPS (64K TPS with 2 trigger points)
- **NRF:** 15K MPS (7.5K TPS with 2 trigger points)
- **SEPP:** 60K MPS (15K TPS with 4 trigger points)
- **BSF:** 10K MPS (2 trigger points)
- **PCF:** 58K MPS (2 trigger points)

Benchmark Test Results

Traffic Feed Details

Table 4-11 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	270K MPS	12	33.6ms	Feed: 99.98%

Table 4-11 (Cont.) Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
PCF	Ingress Gateway (Request) Egress Gateway (Request)	58K TPS	12	NA	99.99%
SCP	Ingress Gateway (Request) Egress Gateway (Request)	128K TPS	12	NA	99.95%
BSF	Ingress Gateway (Request) Egress Gateway (Request)	10K TPS	12	NA	100%
NRF	Ingress Gateway (Request) Egress Gateway (Request)	15K TPS	12	NA	99.98%
SEPP	Ingress plmn/n32 Gateway (Request) Egress plmn/n32 Gateway (Request)	60K TPS	12	NA	99.91%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-12 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
ocnaddnrfaggregation	62.3	35.5	1
ocnaddpcfaggregation	61.8	30.7	2
ocnaddbsfaggregation	40.3	17.9	1
ocnaddscpaggregation	65.3	21.3	5
ocnaddseppaggregation	60.2	18.7	3
adapter	42.6	55.9	23
kafka-broker	61.1	42.6	11

Table 4-12 (Cont.) OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
kraft-controller	0.448	50.3	3
ocnaddadminservice	0.0581	28.5	1
ocnaddalarm	5.77	32.3	1
ocnaddconfiguration	0.135	34.1	1
ocnaddgui	0.000994	3.32	1
ocnaddhealthmonitoring	0.496	33.6	1
ocnadduirouter	0.00817	29.7	1

4.1.5 Performance Benchmarking for 270K MPS Aggregated Traffic with Synthetic (TCP) Feed

This performance benchmarking evaluates the deployment of OCNADD handling an aggregated feed of 270K MPS. The breakdown includes SCP traffic (64K TPS with 2 trigger points), NRF 15K MPS (7.5K TPS with 2 trigger points) and SEPP 60K MPS (15K TPS with 4 trigger points), BSF 10K MPS (with 2 trigger points) PCF 58K MPS (with 2 trigger points). The benchmark utilizes a synthetic (TCP) feed:

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

Execution Details

- **Execution Time:** 12+ Hours

Call Rate

- **SCP Traffic:** 128K MPS (64K TPS with 2 trigger points)
- **NRF:** 15K MPS (7.5K TPS with 2 trigger points)
- **SEPP:** 60K MPS (15K TPS with 4 trigger points)
- **BSF:** 10K MPS (2 trigger points)
- **PCF:** 58K MPS (2 trigger points)

Call Mix

- **NF Call Mix:** SCP, NRF, SEPP, BSF and PCF NF Call mix models

Setup Details

- **Environment:** OCCNE 24.3.1
- **LB Type:** Bare metal CNLB
- **Cluster Topology:**

- **DD:** Cluster-1
- **NF:** Cluster-2
- **Third-Party Consumers:** Cluster-3
- **OCNADD Version:** 25.1.200
- **cnDBTier:** 25.1.100

Network Functions (NFs)

- **SCP:** 25.1.100
- **NRF:** 25.1.100
- **SEPP:** 25.1.100
- **BSF:** 25.1.100
- **PCF:** 25.1.100

Third-Party Consumers

- Two endpoints with two different egress NADs (round-robin)

Configuration Details

- **NRF:** Ingress Gateway and Egress Gateway `messageCopy` enabled
- **SCP:** Default configuration
- **cSEPP:**Ingress PLMN/N32 and Egress PLMN/N32 `messageCopy` enabled
- **pSEPP:**`messageCopy` disabled
- **PCF:** Ingress Gateway and Egress Gateway `messageCopy` enabled
- **BSF:** Ingress Gateway and Egress Gateway `messageCopy` enabled

Security and Data Feed

- **SASL/SSL:** Enabled between BSF and OCNADD
- **OCNADD Replication Factor:** 1
- **OCNADD Data Feed:** HTTP2
- **OCNADD Kafka PVC:** 400 GB
- **Average Message Size:** ~3500 Bytes

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with aggregation rules set for SCP, NRF, SEPP, BSF, and PCF.
 - Message Ingestion Rate: 270K MPS

Resource Specifications:

Table 4-13 OCNADD Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddscppaggregation	2	2	4	4	5	5	30	SCP
ocnaddnrffaggregation	2	2	2	2	1	1	6	NRF
ocnaddseppaggregation	2	2	4	4	3	3	18	SEPP
ocnaddadapter	3	3	6	6	15	15	90	MAIN
ocnaddkafka	4	4	48	48	11	11	-	-
kraft-controller	1	1	1	2	3	3	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnadduirouter	1	2	1	1	1	2	-	-
ocnaddfilter	2	2	3	3	1	1	-	-
ocnaddexport	2	4	4	64	1	2	-	-
ocnaddpcfaggregation	2	2	2	2	3	3	18	PCF
ocnaddbsfaggregation	2	2	2	2	1	1	6	BSF

- **SCP Traffic:** 128K MPS (64K TPS with 2 trigger points)
- **NRF:** 15K MPS (7.5K TPS with 2 trigger points)
- **SEPP:** 60K MPS (15K TPS with 4 trigger points)
- **BSF:** 10K MPS (2 trigger points)
- **PCF:** 58K MPS (2 trigger points)

Benchmark Test Results

Traffic Feed Details

Table 4-14 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	270K MPS	12	33.6ms	Feed: 99.98%

Table 4-14 (Cont.) Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
PCF	Ingress Gateway (Request) Egress Gateway (Request)	58K TPS	12	NA	99.99%
SCP	Ingress Gateway (Request) Egress Gateway (Request)	128K TPS	12	NA	99.95%
BSF	Ingress Gateway (Request) Egress Gateway (Request)	10K TPS	12	NA	100%
NRF	Ingress Gateway (Request) Egress Gateway (Request)	15K TPS	12	NA	99.98%
SEPP	Ingress plmn/n32 Gateway (Request) Egress plmn/n32 Gateway (Request)	60K TPS	12	NA	99.91%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-15 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
ocnaddnrfaggregation	58.4	31.9	1
ocnaddpcfaggregation	55.2	27.8	3
ocnaddbsfaggregation	35.4	20.9	1
ocnaddscpaggregation	62.4	19.9	5
ocnaddseppaggregation	58.2	17.1	3
adapter	59.5	13.1	15
kafka-broker	57.2	11.6	11

Table 4-15 (Cont.) OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
kraft-controller	0.455	42.4	3
ocnaddadminservice	0.417	29.3	1
ocnaddalarm	0.137	33.7	1
ocnaddconfiguration	0.129	34.2	1
ocnaddgui	0.00100	3.32	1
ocnaddhealthmonitoring	0.382	33.8	1
ocnadduirouter	0.00838	29.5	1

4.1.6 Performance Benchmarking for 270K MPS SCP Traffic with Kafka Feed

This performance benchmarking evaluates a centralized deployment with a single worker group, handling an SCP traffic of 270K MPS. The benchmark utilizes a Kafka feed:

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

- **Ingress Traffic Rate:** 270K MPS
- **Observation Period:** 3+ hours
- **Deployment:** OCNADD single-site, ASM disabled
- **Kafka Replication Factor:** 1
- **Kafka PVC:** 300GB
- **Data Feed:** Kafka-feed
- **Third-Party Consumers:** 1 endpoint per feed
- **Software Versions:**
 - **cnDBTier:** 24.3.0
 - **SCP:** 24.3.0
 - **CNCC:** 24.3.0
- **Security:** SASL/SSL enabled between between SCP and OCNADD
- **Message Size:** ~3500 - 4000 Bytes
- **Execution Time:** 12 hours
- **Call Mix:** SCP SCP Model-C

Setup Details

- **Environment:** OCCNE 23.3.4

- **LB Type:** Bare metal LBVM
- **OCNADD:** 25.1.100

Configuration

- **SCP:** Default configuration, messageCopy enabled
- **OCNADD:**
 - Replication Factor: 1
 - Kafka PVC: 300GB. For more information about PVC requirements, see [Kafka PVC Storage Requirements](#).

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with SCP.
 - Message Ingestion Rate: 270 K MPS

Resource Specifications:

Table 4-16 Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Memory Request Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddadapter	-	-	-	-	-	-	261	MAIN
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddexport	2	4	4	64	1	2	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddkafka	5	5	48	48	11	11	-	-
ocnaddscppaggregation	2	2	4	4	11	11	66	SCP
ocnadduirouter	1	2	1	1	1	2	-	-
zookeeper	1	1	1	2	3	3	-	-

- **SCP:** SCP Traffic is 135K TPS using two trigger points.

Benchmark Test Results

Traffic Feed Details

Table 4-17 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	270K MPS	12	Feed-1: 26.9 msec Feed-2: 26.5 msec	Feed-1: 99.98% Feed-2: 99.98%
SCP	Ingress Gateway (Request) Egress Gateway (Request)	135K TPS	12	NA	99.98%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-18 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
kafka-broker	46.1	16.3	11
ocnaddadminservice	0.0855	25.9	1
ocnaddalarm	0.80	26.8	1
ocnaddconfiguration	0.139	33.2	1
ocnaddexport	0.0115	2.18	1
ocnaddgui	0.000999	3.10	1
ocnaddhealthmonitoring	0.339	35.0	1
ocnaddscpaggregation	68.2	27.9	11
ocnadduirouter	0.00654	54.0	1
zookeeper	0.121	28.8	3

4.1.7 Performance Benchmarking for 270K MPS SCP Traffic with Synthetic Feed

This performance benchmarking evaluates a centralized deployment with a single worker group, handling an SCP Model-C traffic of 270K MPS. The benchmark utilizes a synthetic feed:

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

- Ingress Traffic Rate:** 270K MPS

- **Observation Period:** 3+ hours
- **Deployment:** OCNADD single-site, ASM disabled
- **Kafka Replication Factor:** 1
- **Kafka PVC:** 300GB
- **Data Feed:** TCP (Synthetic)
- **Third-Party Applications:** 1 endpoint per feed
- **Software Versions:**
 - **cnDBTier:** 24.3.0
 - **NRF:** 24.3.0
 - **CNCC:** 24.3.0
- **Security:** SASL/SSL enabled between SCP and OCNADD
- **Message Size:** ~3500 - 4000 Bytes
- **Execution Time:** 14 hours
- **Call Mix:** SCP Model-C

Setup Details

- **Environment:** OCCNE 23.3.4
- **LB Type:** Bare metal LBVM
- **OCNADD:** 25.1.100

Configuration

- **SCP:** Default configuration, messageCopy enabled
- **OCNADD:**
 - Replication Factor: 1
 - Kafka PVC: 300GB. For more information about PVC requirements, see [Kafka PVC Storage Requirements](#).

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with SCP.
 - Message Ingestion Rate: 270K MPS

Resource Specifications:

Table 4-19 OCNADD Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddadapter(TCPs)	3	3	6	6	21	21	126	MAIN

Table 4-19 (Cont.) OCNADD Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddexport	2	4	4	64	1	2	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddkafka	5	5	48	48	11	11	-	-
ocnaddscppaggregation	2	2	4	4	11	11	66	SCP
ocnadduirouter	1	2	1	1	1	2	-	-
zookeeper	1	1	1	2	3	3	-	-

- **SCP:** SCP Traffic is 135K TPS using two trigger points.

Benchmark Test Results

Traffic Feed Details

Table 4-20 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	270K MPS	14	51.1 msec	99.99%
SCP	Ingress Gateway (Request) Egress Gateway (Request)	135K MPS	14	NA	99.98%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-21 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
ocnaddscppaggregation	65.5	29.8	11
ocnaddadapter(TCPs)	58.1	29.0	21
kafka-broker	44.1	60.5	11
zookeeper	0.152	33.7	3

Table 4-21 (Cont.) OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
ocnaddadminservice	0.134	27.9	1
ocnaddalarm	0.399	18.4	1
ocnaddconfiguration	0.232	33.9	1
ocnaddexport	0.0183	2.21	1
ocnaddgui	0.00110	3.12	1
ocnaddhealthmonitoring	0.710	33.2	1
ocnadduirouter	0.00812	25.3	1

4.1.8 Performance Benchmarking for 135K MPS Traffic with Synthetic Feed Replication

This performance benchmarking evaluates a centralized deployment with a single worker group, handling an aggregated feed of 135K MPS. The breakdown includes 90K MPS for SCP (45K TPS with 2 trigger points), 30K MPS for SEPP (15K TPS with 2 trigger points), and 15K MPS for NRF (7.5K TPS with 2 trigger points). The benchmark utilizes a replicated synthetic feed:

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

- **Ingress Traffic Rate:** 135K MPS
- **Observation Period:** 3+ hours
- **Deployment:** OCNADD single-site, ASM disabled
- **Kafka Replication Factor:** 1
- **Kafka PVC:** 400GB
- **Data Feed:** Two TCP (Synthetic) Feeds (TLS enabled)
- **Third-Party Applications:** 1 endpoint per feed
- **Software Versions:**
 - **cnDBTier:** 24.3.0
 - **SCP:** 24.3.0
 - **NRF:** 24.3.0
 - **SEPP:** 24.3.0
 - **CNCC:** 24.3.0
- **Security:** SASL/SSL enabled between NRF, SCP, SEPP, and OCNADD
- **Message Size:** ~3500 Bytes

- **Execution Time:** 24 hours
- **Call Mix:** SCP, NRF, and SEPP NF CALL MIX models

Setup Details

- **Environment:** OCCNE 23.3.4
- **LB Type:** Bare metal LBVM
- **OCNADD:** 25.1.100

Configuration

- **NRF:** Ingress Gateway and Egress Gateway messageCopy enabled
- **SCP:** Default configuration
- **cSEPP:** PLMN Ingress Gateway messageCopy enabled
- **pSEPP:** messageCopy disabled
- **OCNADD:**
 - Replication Factor: 1
 - Kafka PVC: 400GB. For more information about PVC requirements, see [Kafka PVC Storage Requirements](#).

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with aggregation rules set for SCP, NRF and SEPP.
 - Message Ingestion Rate: 135K MPS

Resource Specifications:

Table 4-22 OCNADD Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddadapter	3	3	6	6	11	11	66	MAIN
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddexport	2	4	4	64	1	2	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddkafka	7.5	7.5	96	96	6	6	-	-
ocnaddnrfaggregation	2	2	2	2	1	1	6	NRF

Table 4-22 (Cont.) OCNADD Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddscppaggregation	2	2	4	4	5	5	30	SCP
ocnaddseppaggregation	2	2	4	4	2	2	12	SEPP
ocnadduirouter	1	2	1	1	1	2	-	-
zookeeper	1	1	1	2	3	3	-	-

- **SCP:** 90K MPS traffic (45K TPS with 2 trigger points)
- **NRF:** 15K MPS traffic (7.5K TPS with 2 trigger points)
- **SEPP** 30K MPS traffic (15K TPS with 2 trigger points)

Benchmark Test Results

Traffic Feed Details

Table 4-23 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	135K MPS	24	Feed-1: 26.9 msec Feed-2: 26.5 msec	Feed-1: 99.98% Feed-2: 99.98%
SCP	Ingress Gateway (Request) Egress Gateway (Request)	45K TPS	24	NA	99.98%
NRF	Ingress Gateway (Request) Egress Gateway (Request)	7.5K TPS	24	NA	99.98%
SEPP	Ingress Gateway (Request) Egress Gateway (Request)	15K TPS	24	NA	99.98%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-24 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
adapter-1	60.6	19.0	11
adapter-2	63.0	19.4	11
kafka-broker	40.1	25.8	6
ocnaddadminservice	0.138	24.1	1
ocnaddalarm	0.217	27.5	1
ocnaddconfiguration	0.164	34.9	1
ocnaddexport	0.0115	2.33	1
ocnaddgui	0.00108	3.12	1
ocnaddhealthmonitoring	0.530	34.5	1
ocnaddnrffaggregation	70.4	32.4	1
ocnaddscppaggregation	58.0	18.6	5
ocnaddseppaggregation	70.0	20.3	2
ocnadduirouter	0.00758	28.7	1
zookeeper	0.125	31.0	3

4.1.9 Performance Benchmarking for 30K MPS PCF Traffic with HTTP2 Feed

The performance benchmarking test is performed on OCNADD with 31K MPS traffic with HTTP2 feed. The traffic includes:

- PCF Traffic: 30K MPS
- Enabled PCF Ingress Gateway and Egress Gateway Message Copy Feature.

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

- **Ingress Traffic Rate:** 30K MPS
- **Observation Period:** 3+ hours
- **Deployment:** OCNADD single-site, ASM disabled
- **Kafka Replication Factor:** 1
- **Kafka PVC:** 50GB
- **Data Feed:** HTTP2
- **Third-Party Applications:** 1 endpoint per feed
- **Software Versions:**
 - **cnDBTier:** 24.3.0

- **PCF:** 24.3.0
- **CNCC:** 24.3.0
- **Security:** SASL/SSL enabled between PCF and OCNADD
- **Message Size:** ~2079 Bytes
- **Execution Time:** 12 hours
- **Call Mix:** Session Management (SM) traffic

Setup Details

- **Environment:** OCCNE 23.3.4
- **LB Type:** Bare metal LBVM
- **OCNADD:** 25.1.100

Configuration

- **PCF:** Ingress Gateway and Egress Gateway messageCopy enabled
- **OCNADD:**
 - Replication Factor: 1
 - Kafka PVC: 50GB. For more information about PVC requirements, see [Kafka PVC Storage Requirements](#).

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with PCF.
 - Message Ingestion Rate: 30K MPS

Resource Specifications:

Table 4-25 OCNADD Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddadapter	3	3	4	4	3	3	27	MAIN
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddpcfaggregation	2	2	2	2	2	2	12	PCF
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddexport	2	4	4	64	1	2	-	-
ocnaddfilter	2	2	3	3	1	1	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-

Table 4-25 (Cont.) OCNADD Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddkafka	2	2	24	24	4	4	-	-
ocnadduirouter	1	2	1	1	1	2	-	-
zookeeper	1	1	1	2	3	3	-	-

- **PCF** : PCF Traffic: 15K TPS using two trigger point.

Benchmark Test Results

Traffic Feed Details

Table 4-26 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	30K MPS	12	33.6ms	Feed: 99.99%
PCF	Ingress Gateway (Request) Egress Gateway (Request)	15K TPS	12	NA	99.99%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-27 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
httpsfeed-adapter	51.6	62.10	3
kafka-broker	50.8	18.10	4
ocnaddadminservice	0.08	33.90	1
ocnaddalarm	0.12	19.30	1
ocnaddconfiguration	0.09	40.70	1
ocnaddexport	0.03	2.05	1
ocnaddgui	0.00	3.09	1
ocnaddhealthmonitoring	0.13	35.70	1
ocnaddpcfaggregation	61.4	32.90	2
ocnadduirouter	0.01	32.50	1
zookeeper	0.14	48.20	3

4.1.10 Performance Benchmarking for 9K BSF Traffic with HTTP2 Feed

The performance benchmarking test is performed on OCNADD with 9K traffic with HTTPS feed. The traffic includes:

- BSF Traffic: 9K MPS
- Enabled BSF Ingress Gateway and Egress Gateway Message Copy Feature.

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

- **Ingress Traffic Rate:** 9K MPS
- **Observation Period:** 3+ hours
- **Deployment:** OCNADD single-site, ASM disabled
- **Kafka Replication Factor:** 1
- **Kafka PVC:** 30GB
- **Data Feed:** HTTP2
- **Third-Party Applications:** 1 endpoint per feed
- **Software Versions:**
 - **cnDBTier:** 24.3.0
 - **BSF:** 24.3.0
 - **CNCC:** 24.3.0
- **Security:** SASL/SSL enabled between between BSF and OCNADD
- **Message Size:** ~1200 Bytes
- **Execution Time:** 12 hours
- **Call Mix:** Binding-Create/Delete

Setup Details

- **Environment:** OCCNE 23.3.4
- **LB Type:** Bare metal LBVM
- **OCNADD:** 25.1.100

Configuration

- **BSF:** Ingress Gateway and Egress Gateway messageCopy enabled
- **OCNADD:**
 - Replication Factor: 1
 - Kafka PVC: 30GB. For more information about PVC requirements, see [Kafka PVC Storage Requirements](#).

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with BSF.
 - Message Ingestion Rate: 9K MPS

Resource Specifications:

Table 4-28 OCNADD Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddadapter	3	3	4	4	1	1	9	MAIN
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddbsfaggregation	3	3	4	4	1	1	6	BSF
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddexport	2	4	4	64	1	2	-	-
ocnaddfilter	2	2	3	3	1	1	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddkafka	2	2	24	24	4	4	-	-
ocnaddrouter	1	2	1	1	1	2	-	-
zookeeper	1	1	1	2	3	3	-	-

- **BSF** : BSF Traffic: 4.5K TPS using two trigger points.

Benchmark Test Results**Traffic Feed Details****Table 4-29 Traffic Feed Details**

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	9K MPS	12	20.20ms	Feed: 100%
BSF	Ingress Gateway (Request) Egress Gateway (Request)	4.5K TPS	12	NA	100%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-30 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
adapter	33.5	32.1	1
kafka-broker	12.3	11.4	4
ocnaddadminservice	0.287	27.5	1
ocnaddalarm	0.0963	37.7	1
ocnaddbsfaggregation	31.0	22.1	1
ocnaddconfiguration	0.110	36.7	1
ocnaddexport	0.00686	2.32	1
ocnaddgui	0.00107	3.13	1
ocnaddhealthmonitoring	0.0898	35.1	1
ocnadduirouter	0.00664	29.5	1
zookeeper	0.122	27.0	3

4.2 vCNE Cluster Environment

This section gives the details of the performance tests done in LBVM vCNE Cluster Environment.

4.2.1 Performance Benchmarking for 360K MPS SCP and 450K Egress Traffic with TCP and HTTP2 Feed

This performance benchmarking evaluates the deployment of OCNADD handling an SCP traffic of 360K MPS and Egress traffic of 450K MPS. The benchmark utilizes 3 TCP feeds and 1 HTTP2 feed:

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

Execution Details

- **Execution Time:** 12+ Hours

Call Rate

- **Total Traffic:** 360K MPS

Call Mix

- **Model:** SCP Model-C

Setup Details

- **Environment:** OCCNE 24.3.1
- **LB Type:** vCNE LBVM
- **OCNADD Version:** 25.1.200
- **cnDBTier:** 24.3.0
- **NF Version:** 25.1.100
- **Third-Party Application:** One endpoint configured

Configuration Details

- **SCP:** Worker messageCopy enabled
- **SASL/SSL:** Enabled between SCP and OCNADD
- **OCNADD Replication Factor:** 1
- **OCNADD Data Feed:** HTTP2
- **OCNADD Kafka PVC:** 400 Gi
- **Average Message Size:** 3564 Bytes

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with SCP.
 - Message Ingestion Rate: 360K MPS

Resource Specifications:

Table 4-31 Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddscppaggregation	2	2	4	4	13	13	78	SCP
tcpfeed-adapter	3	3	6	6	20	20	120	MAIN
tcpfeed-2-adapter	3	3	10	10	7	7	120	MAIN
tcpfeed-3-adapter	3	3	10	10	7	7	120	MAIN
http2feed-adapter	3	3	28	28	5	5	120	MAIN
ocnaddkafka	5	5	96	96	16	16	-	-
kraft-controller	1	1	2	2	3	3	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnadduirouter	1	2	1	1	1	2	-	-

Table 4-31 (Cont.) Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddfilter	2	2	3	3	1	1	-	-
ocnaddexport	2	4	4	64	1	2	-	-

- **SCP:** SCP Traffic is 180K TPS using two trigger points.

Benchmark Test Results

Traffic Feed Details

Table 4-32 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	360K MPS	12	http: 8ms tcp(360k): 20ms tcp(18k): 10ms	Feed: 99.99%
SCP	Worker	180K TPS	12	NA	99.99%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-33 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
ocnaddscpaggregation	77	20	13
tcpfeed-adapter	75.40	38.97	20
tcpfeed-2-adapter	81	31	7
tcpfeed-3-adapter	85	21	7
http2feed-adapter	80.60	58.52	5
kafka-broker	31	10.6	16
kraft-controller	3	51	3
zookeeper	0	0	0
ocnaddadminservice	0.20	36.80	1
ocnaddalarm	0.20	45.10	1
ocnaddconfiguration	0.10	47.60	1
ocnaddexport	0.10	56.90	1
ocnaddgui	0.10	3.30	1
ocnaddhealthmonitoring	0.20	38.80	1
ocnadduirouter	0.10	34.80	1

4.2.2 Performance Benchmarking for 360K MPS SCP Traffic with Synthetic (TCP) Feed

This performance benchmarking evaluates the deployment of OCNADD handling an SCP traffic of 360K MPS. The benchmark utilizes a TCP feed:

The benchmarking results are provided below:

Note

One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

Execution Details

- **Execution Time:** 12+ Hours

Call Rate

- **Total Traffic:** 360K MPS

Call Mix

- **Model:** SCP Model-C

Setup Details

- **Environment:** OCCNE 24.3.1
- **LB Type:** vCNE LBVM
- **OCNADD Version:** 25.1.200
- **cnDBTier:** 25.1.100
- **SCP Version:** 25.1.100
- **Third-Party Application:** One endpoint configured

Configuration Details

- **SCP:** Default configuration

Security & Data Feed

- **OCNADD Replication Factor:** 1
- **OCNADD Data Feed:** Synthetic (TCP)
- **OCNADD Kafka PVC:** 400 GB
- **Average Message Size:** ~3500 Bytes

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with SCP.
 - Message Ingestion Rate: 360K MPS

Resource Specifications:

Table 4-34 OCNADD Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Memory Request Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Partitions/Retention	Topics
ocnaddconfiguration	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddhealthmonitoring	1	1	1	1	1	1	-	-
ocnaddscppaggregation	2	2	4	4	13	13	78	SCP
ocnaddadapter	3	3	6	6	20	20	120	MAIN
ocnaddkafka	5	5	48	48	11	11	-	-
kraft-controller	1	1	2	2	3	3	-	-
ocnaddgui	1	2	1	1	1	2	-	-
ocnadduirouter	1	2	1	1	1	2	-	-

- **SCP:** SCP Traffic is 180K TPS using two trigger points

Benchmark Test Results**Traffic Feed Details****Table 4-35 Traffic Feed Details**

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	360K MPS	12	22.6ms	Feed: 99.99%
SCP	Ingress Gateway (Request) Egress Gateway (Request)	180K TPS	12	NA	99.99%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-36 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
ocnaddscppaggregation	77	20	13

Table 4-36 (Cont.) OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
TCP: adapter (360k)	78.4	36.97	20
kafka-broker	35	15	11
kraft-controller	5	48	3
ocnaddadminservice	0.2	36.8	1
ocnaddalarm	0.2	45.1	1
ocnaddconfiguration	0.1	47.6	1
ocnaddgui	0.1	3.3	1
ocnaddhealthmonitoring	0.2	38.8	1
ocnaddrouter	0.1	34.8	1